Open in app ↗ Sign up Sign in

Medium Search

The need for a Thinking, Private Brain & The Death of Dumb Drones

20 min read · Just now

Berend Watchus Follow

Listen Share

Author: Berend Watchus December 10, 2025

Observing, Orienting, Deciding … locally and privately

Figure 1: Local OODA loops in a coordinated UAV swarm. Each drone maintains its own Observe-Orient-

Decide-Act cycle (white arrows) while communicating securely with the swarm network (cyan links). This distributed intelligence architecture enables real-time autonomous coordination without centralized control — the foundation of privacy-preserving LLM-driven swarms. **Note: This is a conceptual illustration for educational purposes; actual communication topology, data flow patterns, and network density in deployed systems vary based on mission parameters and are not represented at scale.**

The threat landscape has fundamentally changed. It's no longer enough to encrypt a video stream and hope for the best. We are exiting the era of the Remote-Controlled (RC) drone and entering the age of the **Autonomous Intelligence Agent.**

3,700× Faster: The Tiny Chip That Makes Quantum-Safe IoT Affordable

Author: Berend Watchus + 'Autonomous Knowledge Accelerator'

medium.com

EXCLUSIVE: The Autonomous Knowledge Accelerator — How a Commercial LLM Generated 200× Power and...

Author: Berend Watchus

systemweakness.com

My recent research — from the **Quantum Shield Chip (QSC-1)** to the **Autonomous Knowledge Accelerator** — shows that the technology exists to build systems that are hundreds of times faster and more efficient. Now, a new paper on the arXiv, **'PrivLLMSwarm: Privacy-Preserving LLM-Driven UAV Swarms for Secure IoT Surveillance'**, provides the crucial **operational blueprint** for how to do it securely.

PrivLLMSwarm: Privacy-Preserving LLM-Driven UAV Swarms for Secure IoT Surveillance

Large Language Models (LLMs) are emerging as powerful enablers for autonomous reasoning and natural-language...

arxiv.org

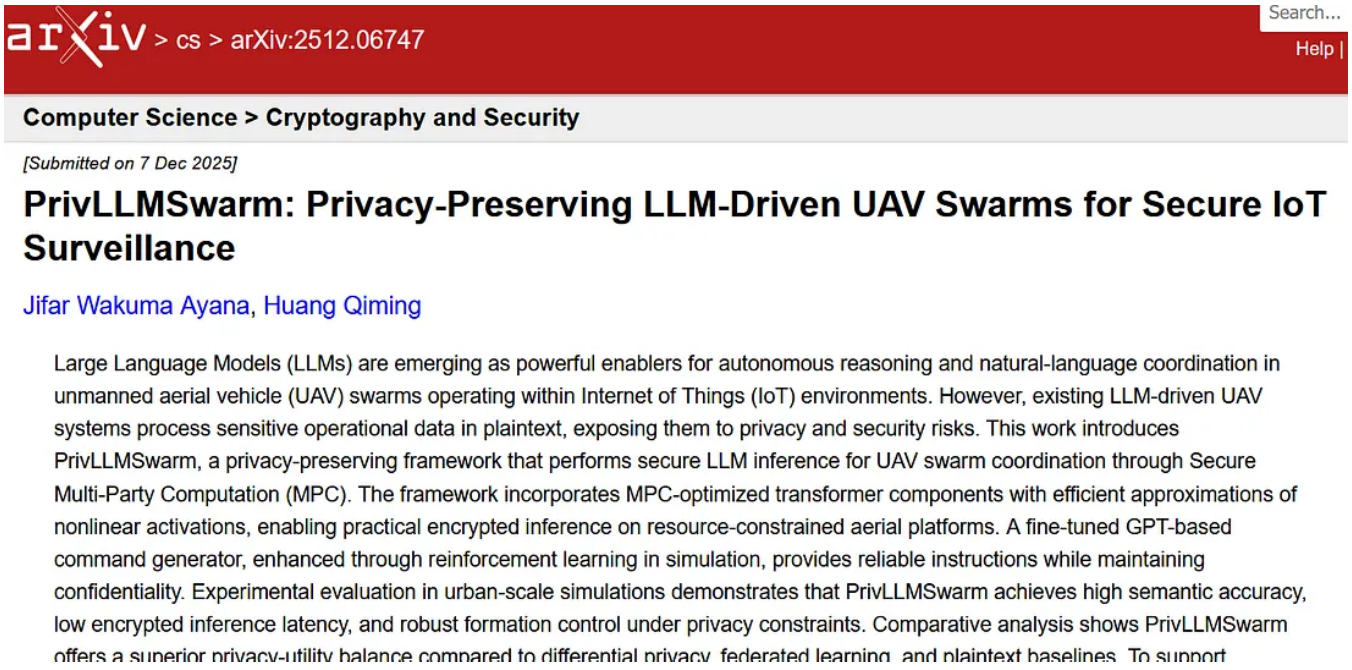

arXiv > cs > arXiv:2512.06747

Search... Help |

Computer Science > Cryptography and Security

[Submitted on 7 Dec 2025]

PrivLLMSwarm: Privacy-Preserving LLM-Driven UAV Swarms for Secure IoT Surveillance

Jifar Wakuma Ayana, Huang Qiming

Large Language Models (LLMs) are emerging as powerful enablers for autonomous reasoning and natural-language coordination in unmanned aerial vehicle (UAV) swarms operating within Internet of Things (IoT) environments. However, existing LLM-driven UAV systems process sensitive operational data in plaintext, exposing them to privacy and security risks. This work introduces PrivLLMSwarm, a privacy-preserving framework that performs secure LLM inference for UAV swarm coordination through Secure Multi-Party Computation (MPC). The framework incorporates MPC-optimized transformer components with efficient approximations of nonlinear activations, enabling practical encrypted inference on resource-constrained aerial platforms. A fine-tuned GPT-based command generator, enhanced through reinforcement learning in simulation, provides reliable instructions while maintaining confidentiality. Experimental evaluation in urban-scale simulations demonstrates that PrivLLMSwarm achieves high semantic accuracy, low encrypted inference latency, and robust formation control under privacy constraints. Comparative analysis shows PrivLLMSwarm offers a superior privacy-utility balance compared to differential privacy, federated learning, and plaintext baselines. To support

https://arxiv.org/abs/2512.06747 copyright: arXiv and Jifar Wakuma Ayana and Huang Qiming

This is the strategic warning you need to heed: **The race for strategic advantage will be won by the swarm that can think, and think in private.**

1. The Revolution: From “Go North” to On-Board Reasoning

For years, the drone (UAV) swarm model was inherently slow and centralized:

> **The Old Model: Encrypt and Transmit.** A surveillance drone would collect a video stream or sensor data, encrypt it, and send the massive data packet back to a distant command center. The drone’s on-board brain was capable of little more than “Maintain altitude” or “Go North.” The heavy lifting — the crucial analysis and decision-making — was handled in one of two ways:

1. **For slow, strategic intelligence gathering:** It happened hours later in a secure, human-operated base.
2. **For real-time tactical or logistics missions:** It happened live at that moment by a **human operator** who received the encrypted stream and then manually issued

simple commands like ‘go north’ or ‘land’.

> Both approaches create a catastrophic failure in terms of **urgency** and **cost**. They rely on a fragile communication link, waste bandwidth, burn power, and, worst of all, introduce human reaction latency into a machine-speed environment.

The Novelty of PrivLLMSwarm: The Thinking Swarm

The *PrivLLMSwarm* paper flips this model on its head. The novelty is the integration of an LLM (Large Language Model) — the analytical brain — directly onto the swarm agents, while simultaneously keeping the entire operation confidential.

- **LLM-Driven Coordination:** The drone is no longer a sensor platform; it becomes an **autonomous reasoning entity**. It can process diverse inputs (visual, textual, sensor data) and generate context-aware, complex commands like: “Alpha-Drone 3: Shift formation to wedge pattern, reduce altitude by 20 meters, and investigate the secondary heat signature near the water tower.”
- **Secure Multi-Party Computation (MPC):** This is the game-changer. The LLM’s inference (its “thinking”) and the sensitive data it uses are protected using **MPC**. The inference happens while the data remains mathematically encrypted, meaning no single drone, nor any eavesdropper, can read the plaintext data or the proprietary model weights.

The Significance: *PrivLLMSwarm* establishes a **practical foundation** for secure, real-time coordination that balances privacy (critical for dual-use applications like smart-city monitoring) and operational effectiveness (critical for battlefield or disaster scenarios).

2. The Core Mechanism: The Constant Verbal OODA Loop

The true power of an LLM-driven swarm is its ability to constantly and rapidly cycle through the **OODA Loop** (Observe, Orient, Decide, Act).

In a traditional system, the Orient and Decide phases were slow, external, and

human-dependent. In the *PrivLLMSwarm*, the LLM is the internal engine driving this cycle in milliseconds:

- **Observe:** The drone swarm's sensors collect new intel (e.g., visual data, electromagnetic signatures, new commands from the central authority).
- **Orient (The Strategic Brain):** The on-board LLM contextualizes the observation. It doesn't just see a red object; it orients: *"This is an unflagged vehicle in a restricted zone. My **strategic mandate** (mission rules) prioritizes target-of-opportunity over maintaining stealth. However, the current **tactical decision** must prioritize safety due to high wind shear."*
- **Decide:** Based on its orientation, the LLM generates a complex, nuanced action plan in natural language.
- **Act:** The drone executes the complex, coordinated command.

This constant feedback mechanism is the **Verbal OODA Loop**, where the LLM's reasoning is the central, rapid process that translates raw data into coherent, multi-agent actions. The agent is constantly adapting to new intel, new commands, and its own observations, entirely on the edge.

3. The 'Watchus Synthesis': Accelerating and Securing the Brain

While the *PrivLLMSwarm* paper provides the critical software framework (MPC + LLM) for secure thinking, my research provides the **hardware and cognitive architecture** necessary to take this system from an academic proof-of-concept to an undeniable **competitive advantage** in real-world deployment.

Enhancing PrivLLMSwarm with Watchus Architecture

- **QSC-1: The Quantum Shield Chip — Secures the Future & Cuts Cost.** This chip solves the **Urgency** of the quantum threat (Harvest Now, Decrypt Later) by integrating quantum-safe cryptography (PQC) and Zero-Knowledge Proofs (ZKP) into the hardware. Its **$0.45/unit cost** and **3,700× efficiency** makes PQC affordable, directly solving the *PrivLLMSwarm*'s resource-constrained challenge.

The latency of MPC is no longer bottlenecked by slow, unoptimized software.

- **ED-HPSNN / Knowledge Accelerator — 200× Power Efficiency.** The **200× power efficiency** of the Neuromorphic-Photonic architecture (ED-HPSNN) makes the high-power LLM inference on the drone swarm sustainable, drastically extending loiter time. This accelerates the **Observe** phase and reduces operational **cost.**

EXCLUSIVE: The Autonomous Knowledge Accelerator — How a Commercial LLM Generated 200× Power and...

Author: Berend Watchus

systemweakness.com

- **Law of Optimized Complexity (LOOC) — 16.4× Cognitive Speed.** The LLM-driven *Orient* and *Decide* phases, which involve complex reasoning, are guided by the LOOC framework, which delivers a **16.4× improvement in CER.** This ensures the LLM's complex decisions are not only faster but more truthful and reliable, directly enhancing the quality of the LLM's reasoning.

https://www.scribd.com/document/937757981<<

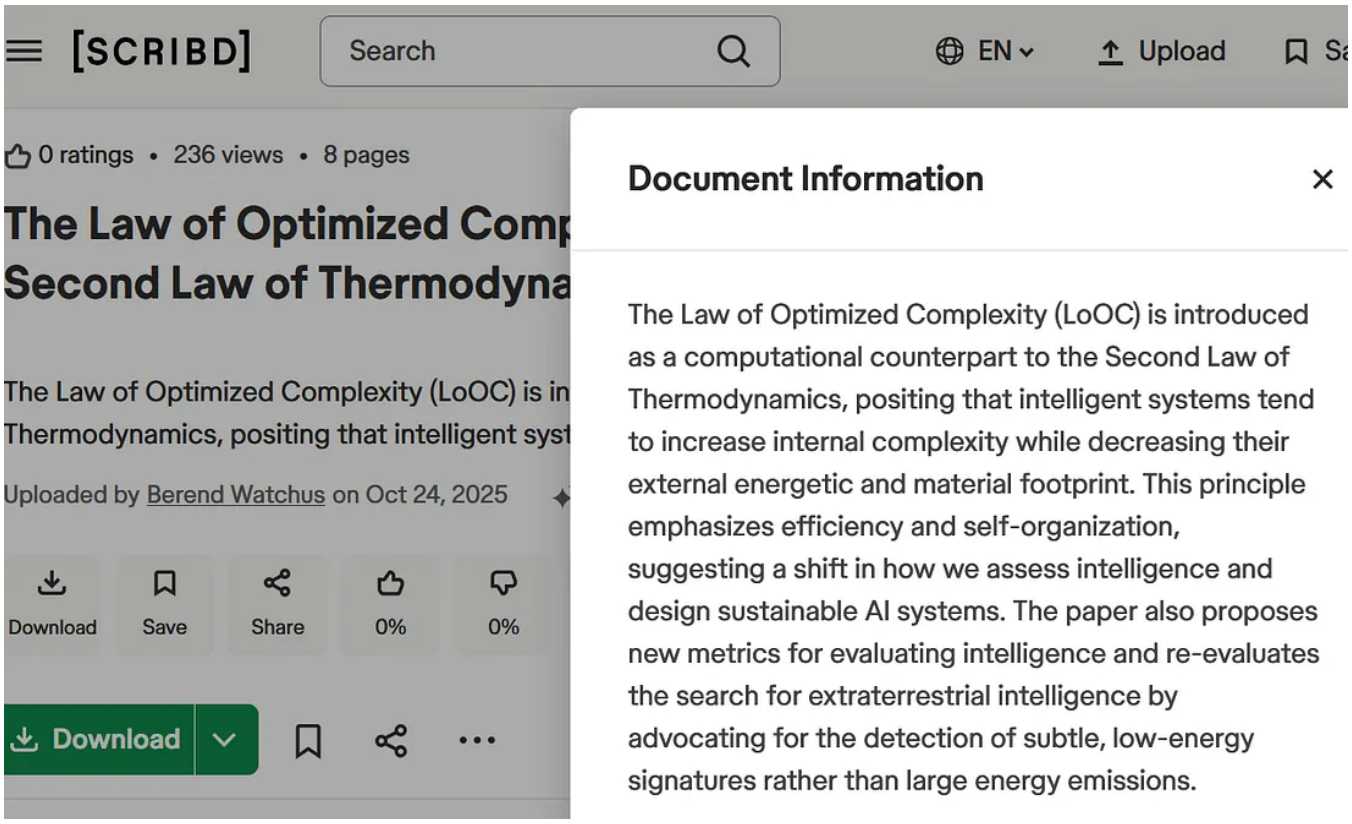
[SCRIBD]
Search
EN
Upload
0 ratings • 236 views • 8 pages
The Law of Optimized Comp
Second Law of Thermodyna
The Law of Optimized Complexity (LoOC) is in
Thermodynamics, positing that intelligent syst
Uploaded by Berend Watchus on Oct 24, 2025
Download
Save
Share
0%
0%
Download
Document Information
The Law of Optimized Complexity (LoOC) is introduced as a computational counterpart to the Second Law of Thermodynamics, positing that intelligent systems tend to increase internal complexity while decreasing their external energetic and material footprint. This principle emphasizes efficiency and self-organization, suggesting a shift in how we assess intelligence and design sustainable AI systems. The paper also proposes new metrics for evaluating intelligence and re-evaluates the search for extraterrestrial intelligence by advocating for the detection of subtle, low-energy signatures rather than large energy emissions.

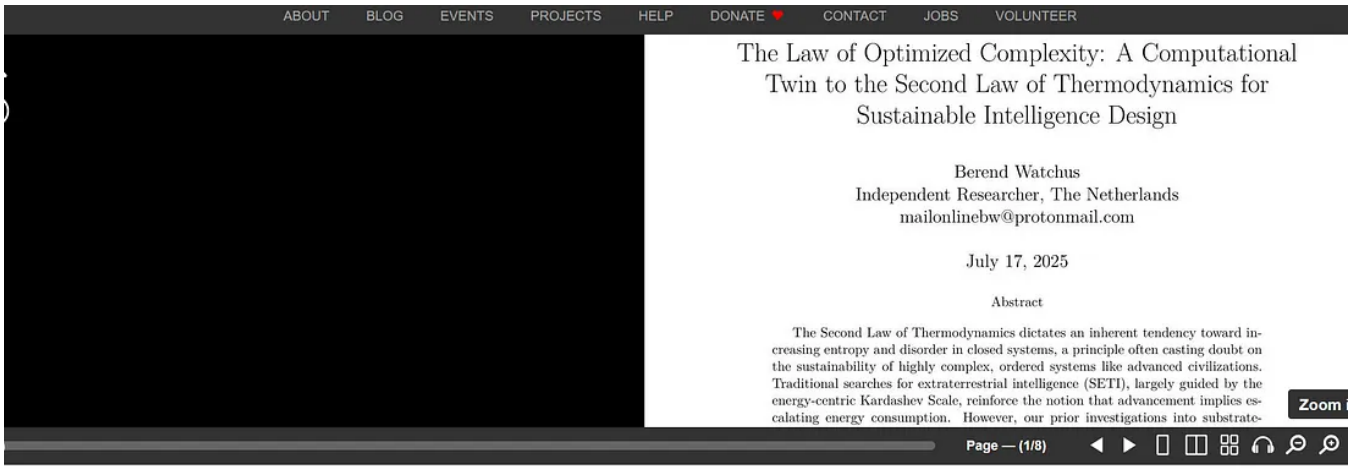
ABOUT BLOG EVENTS PROJECTS HELP DONATE CONTACT JOBS VOLUNTEER
The Law of Optimized Complexity: A Computational Twin to the Second Law of Thermodynamics for Sustainable Intelligence Design
Berend Watchus
Independent Researcher, The Netherlands
mailonlinebw@protonmail.com
July 17, 2025
Abstract
The Second Law of Thermodynamics dictates an inherent tendency toward increasing entropy and disorder in closed systems, a principle often casting doubt on the sustainability of highly complex, ordered systems like advanced civilizations. Traditional searches for extraterrestrial intelligence (SETI), largely guided by the energy-centric Kardashev Scale, reinforce the notion that advancement implies escalating energy consumption. However, our prior investigations into substrate-
Zoom
Page — (1/8)
The Law Of Optimized Complexity A Computational Twin To The Second Law Of Thermodynamics For Sustainable Intelligence Design Berend Watchus
Add to list
Favorite
Share
Flag

https://archive.org/details/the-law-of-optimized-complexity-a-computational-twin-to-the-second-law-of-thermo_202510

https://archive.org/details/the-law-of-optimized-complexity-a-computational-twin-to-the-second-law-of-thermo_202510<<

- **Unified Model of Consciousness — The Adaptive Brain.** This provides the theoretical validation for why the swarm, when equipped with these nested OODA loops and high-speed feedback, exhibits complex, adaptive behavior (like the **Synthetic Insula**). This validates the swarm's capability for genuine autonomous *Orient* functionality.

The Unified Model of Consciousness: Interface and Feedback Loop as the Core of Sentience

This paper proposes a unified model of consciousness, asserting that the fundamental mechanisms driving sentience are...

www.preprints.org

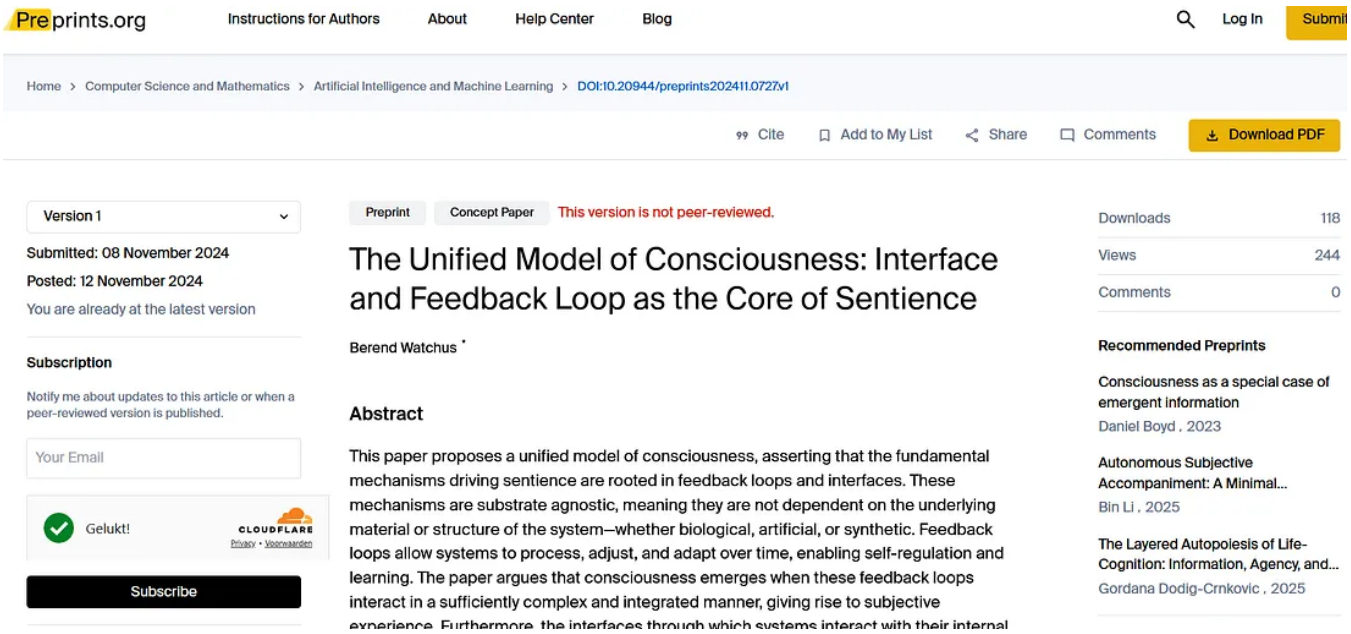
Preprints.org
Instructions for Authors
About
Help Center
Blog
Log In
Submit
Home > Computer Science and Mathematics > Artificial Intelligence and Machine Learning > DOI:10.20944/preprints202411.0727.v1
Cite
Add to My List
Share
Comments
Download PDF
Version 1
Submitted: 08 November 2024
Posted: 12 November 2024
You are already at the latest version
Subscription
Notify me about updates to this article or when a peer-reviewed version is published.
Your Email
Gelukt!
CLOUDFLARE
Privacy • Voorwaarden
Subscribe
Preprint
Concept Paper
This version is not peer-reviewed.
The Unified Model of Consciousness: Interface and Feedback Loop as the Core of Sentience
Berend Watchus *
Abstract
This paper proposes a unified model of consciousness, asserting that the fundamental mechanisms driving sentience are rooted in feedback loops and interfaces. These mechanisms are substrate agnostic, meaning they are not dependent on the underlying material or structure of the system—whether biological, artificial, or synthetic. Feedback loops allow systems to process, adjust, and adapt over time, enabling self-regulation and learning. The paper argues that consciousness emerges when these feedback loops interact in a sufficiently complex and integrated manner, giving rise to subjective
Downloads 118
Views 244
Comments 0
Recommended Preprints
Consciousness as a special case of emergent information
Daniel Boyd , 2023
Autonomous Subjective Accompaniment: A Minimal...
Bin Li , 2025
The Layered Autopoiesis of Life-Cognition: Information, Agency, and...
Gordana Dodig-Crnkovic , 2025

https://www.preprints.org/manuscript/202411.0727

The Integration Mandate

The *PrivLLMSwarm* is the necessary software layer; my work is the enabling hardware and cognitive accelerator.

1. **Integrate QSC-1:** We must replace the reliance on software-only MPC/PQC with the QSC-1's **Hardware-Accelerated MPC-Optimized components** (Rudraksh KEM, zkProof Verifier Core). This takes the "low encrypted inference latency"

achieved by the paper and makes it a field-ready reality that is quantum-immune.

2. **Accelerate the Orient:** We must use the **LOOC-guided, Knowledge-Accelerated LLM** architecture to enhance the LLM used in *PrivLLMSwarm*, maximizing the 16.4× speed gain for the most complex task: translating encrypted data into actionable swarm commands.

This synergy creates a **trustless, quantum-immune, ultra-efficient autonomous swarm** — the ultimate expression of **dual-use** technological dominance for everything from logistics (the U-Hawk concept) to national defense.

Conclusion: The Urgency of the Next Upgrade

The *PrivLLMSwarm* paper should be viewed not just as a breakthrough, but as the first step in a technology path that has just become hyper-accelerated by the **Autonomous Knowledge Accelerator**.

The **urgency** is clear: we must move beyond securing a transmission line and focus on securing the *decision-making process itself*. The **competitive advantage** belongs to the first entity that can deploy an LLM-driven swarm that is both MPC-private *and* QSC-1 quantum-secure.

The next version of the *PrivLLMSwarm* architecture isn't just possible — it's mandatory. It is the synthesis of hardware efficiency, quantum security, and cognitive acceleration.

Disclaimer: Technology as Neutral Infrastructure

Note on Scope Compared to Source Material:

The original PrivLLMSwarm paper (arXiv:2512.06747) focuses exclusively on technical implementation — cryptographic protocols, MPC optimization, and

performance metrics. It does not include ethical analysis, legal considerations, or dual-use discussions beyond mentioning application domains (smart cities, emergency response, infrastructure protection).

This article extends that technical analysis by explicitly addressing the legal, ethical, and strategic implications of deploying secure LLM-driven autonomous systems — a necessary addition for the cybersecurity community and policymakers who must navigate these technologies responsibly.

Note on Technical Novelty:

This article contains zero new inventions, discoveries, or novel technical contributions.

This is an analytical synthesis that:

- Discusses the PrivLLMSwarm paper (arXiv:2512.06747, submitted December 7, 2025)
- References the author's previously published research from 2024–2025, including:
- Quantum Shield Chip (QSC-1) architecture (published November 2025)
- ED-HPSNN neuromorphic computing framework (published October 2025)
- Law of Optimized Complexity (LOOC) (published July 2025)
- Unified Model of Consciousness (published November 2024)
- Provides strategic analysis connecting existing, already-public work

All technical concepts discussed are already in the public domain through:

- Academic preprints (arXiv, Preprints.org)
- Published blog articles (Medium, System Weakness)
- Open-access archives (Archive.org + scribd (was first publsihed with DOI on

zenodo.org)

This article is commentary and synthesis, not original research. It does not:

- ❌ Introduce new cryptographic protocols
- ❌ Present novel hardware designs not previously published
- ❌ Disclose proprietary algorithms or trade secrets
- ❌ Reveal unpublished research results
- ❌ Create new technical vulnerabilities

Purpose: To analyze the convergence of recently published work (PrivLLMSwarm) with the author's existing research portfolio, and to provide the legal and ethical analysis absent from the original technical paper.

This analysis examines the convergence of privacy-preserving AI and autonomous systems as a matter of technological capability and strategic foresight. Like all foundational communications and security infrastructure, these technologies are inherently dual-use and raise significant legal and ethical considerations.

Historical Precedent for Neutral Infrastructure:

The engineers who developed fundamental security protocols created systems that simultaneously protect legitimate users and enable those with malicious intent:

Communications Security:

- WiFi encryption (WPA/WPA2/WPA3)
- Bluetooth security protocols
- GSM/4G/5G cellular encryption

- DECT cordless phone encryption
- End-to-end encrypted messaging (Signal Protocol, used by WhatsApp, Signal, etc.)
- Digital Mobile Radio (DMR) encryption (AES-256)
- P25 Phase 2 public safety radio encryption (AES-256)
- TETRA encrypted trunked radio systems (TEA encryption)
- Modern amateur radio digital modes (D-STAR, DMR, Fusion) with encryption capabilities
- Software-Defined Radio (SDR) platforms with encryption layers

Financial Security Infrastructure:

- Hardware security tokens (RSA SecurID, Yubico keys)
- Banking authentication apps (without physical devices)
- EMV chip card encryption
- SWIFT network message authentication
- TLS/SSL certificates for online banking
- Two-factor authentication systems
- Contactless payment encryption (NFC)

The Amateur Radio Paradox:

Ham radio presents a particularly instructive example. In most jurisdictions, encryption on amateur radio frequencies is legally prohibited to maintain the “open” nature of the hobby and enable regulatory monitoring. Yet the same digital radio technologies used by amateurs (DMR, D-STAR, System Fusion) are commercially deployed with full encryption for:

- Emergency services (police, fire, EMS)
- Private security firms
- Commercial logistics operations
- Military training networks

The technology itself is identical — a Motorola DMR radio or an ICOM D-STAR transceiver can operate in both encrypted and unencrypted modes. The equipment manufacturers don't create "criminal radios" and "legitimate radios" — they create communication infrastructure that serves all users based on configuration and legal authorization.

The Current Reality: Simultaneous Use Right Now

At this very moment — today, December 10, 2025 — the same encrypted radio systems are being used simultaneously by:

Legitimate Users:

- ✅ Paramedics coordinating trauma response in urban emergency departments
- ✅ Firefighters conducting search and rescue in structure fires
- ✅ Police tactical teams executing high-risk warrants
- ✅ Military units in training exercises and active deployments
- ✅ Private security protecting critical infrastructure
- ✅ Disaster response teams coordinating multi-agency relief efforts

Illegitimate Users:

- ⚠️ Drug trafficking organizations coordinating cross-border smuggling operations

- ⚠️ Armed groups in active conflict zones coordinating tactical movements
- ⚠️ Organized crime syndicates managing illegal enterprises
- ⚠️ Human trafficking networks coordinating transportation routes
- ⚠️ Terrorist cells planning and executing attacks

This is happening right now, with the exact same equipment:

- Motorola APX radios (P25 Phase 2, AES-256)
- Hytera DMR portables and mobiles
- Kenwood NEXEDGE digital systems
- Sepura TETRA terminals
- Harris/L3 tactical radios

A paramedic in a major European city and an organized crime operative in another country might both be using Motorola XPR7550 radios with identical encryption capabilities. A police tactical team in one nation and an armed militia group in a conflict zone might both be using Sepura TETRA radios with the same TEA encryption algorithms.

Does this make Motorola, Hytera, Kenwood, or Sepura complicit in organized crime or armed conflict?

No. They manufacture communication infrastructure. The tool is neutral.

The Fundamental Reality:

These systems protect:

- ✅ Legitimate commercial transactions and communications
- ✅ Personal privacy for billions of citizens

- ✅ Military and intelligence operations
- ✅ Corporate intellectual property
- ✅ Financial system integrity
- ✅ Emergency responder safety
- ✅ Critical infrastructure coordination

They also enable:

- ⚠️ Criminal financial operations and communication
- ⚠️ Terrorist communication and funding
- ⚠️ Money laundering and fraud
- ⚠️ Illicit commerce
- ⚠️ Organized crime coordination
- ⚠️ Armed conflict operations
- ⚠️ Human rights abuses

This is not a flaw in the technology — it is the fundamental nature of strong encryption and secure communication systems.

The Supermarket Principle:

A supermarket that sells food does not become complicit in criminal activity simply because criminals also eat. A bank that provides secure transactions does not become responsible for money laundering simply because its encryption works for all customers. The telecommunications company that encrypts phone calls does not endorse the content of those calls. The radio manufacturer that builds encrypted digital transceivers does not control who purchases them or for what purpose.

The infrastructure is neutral; the application determines the ethical and legal

context.

The Real-Time Example:

Right now, as you read this:

- An emergency dispatcher in a major city is using encrypted P25 radio to coordinate an ambulance response to a cardiac arrest
- Simultaneously, somewhere else, encrypted P25 radio is being used to coordinate illegal activities

Both are using the same Motorola APX8000 radio. Same encryption algorithm (AES-256). Same manufacturer. Same technology. Different users, different purposes, different legal frameworks.

The radio manufacturer is not responsible for either use case. The infrastructure is neutral.

Legal Considerations

CRITICAL: Regulations are Dynamic and Jurisdiction-Specific

⚠️ **The legal landscape for dual-use technologies changes frequently and varies dramatically by jurisdiction.**

The European Union regularly updates its **Dual-Use Regulation** (EU 2021/821), expanding or contracting the list of controlled technologies based on emerging capabilities and geopolitical concerns. What is unrestricted today may be export-controlled tomorrow. What requires a license in the EU may be freely exportable in other jurisdictions — or vice versa.

Examples of Regulatory Evolution:

- **2021:** EU adds certain cybersecurity tools to dual-use list
- **2023:** Expanded controls on AI surveillance systems
- **2024:** Enhanced scrutiny of autonomous systems and quantum technologies

- **2025:** Ongoing review of LLM applications with potential security implications
- **Future:** Likely expansion to include privacy-preserving ML and autonomous coordination systems

The United States, China, Russia, India, and other nations maintain separate, evolving export control regimes that may classify identical technologies differently. A cryptographic protocol freely available in one country may be ITAR-restricted or export-controlled in another.

⚠️ **READER RESPONSIBILITY: You are solely responsible for verifying the current legal status of any technology, knowledge, or implementation discussed in this article within your specific jurisdiction at the time you access, use, or deploy it.**

This article provides a snapshot of the regulatory landscape as of December 2025. By the time you read this, regulations may have changed.

Export Controls and Regulatory Frameworks (Current as of December 2025):

The technologies discussed in this article — MPC protocols, post-quantum cryptography, LLM inference systems, and autonomous coordination — are subject to various legal frameworks:

Encryption Export Controls:

- **Wassenaar Arrangement** regulates export of dual-use cryptographic technology (updated periodically by member states)
- **U.S. Export Administration Regulations (EAR)** classify strong encryption as controlled items under Category 5, Part 2
- **EU Dual-Use Regulation (EU 2021/821)** requires licenses for certain cryptographic exports and is regularly amended
- **ITAR (International Traffic in Arms Regulations)** applies to military-grade systems and is subject to State Department revision

Autonomous Systems Regulations:

- **EU AI Act (2024)** classifies autonomous systems as high-risk AI applications — subject to ongoing implementation guidance
- **FAA Part 107** (U.S.) and **EASA regulations** (EU) govern UAV operations and are frequently updated
- **DOD Directive 3000.09** establishes U.S. policy on autonomous weapon systems (subject to revision)
- **Geneva Conventions Additional Protocol I** addresses autonomous targeting concerns (interpretation evolving)

Privacy and Data Protection:

- **GDPR (General Data Protection Regulation)** applies to IoT surveillance data — enforcement priorities shift based on EDPB guidance
- **California Consumer Privacy Act (CCPA)** regulates data collection and is amended annually
- **ePrivacy Directive** covers electronic communications confidentiality — major revision pending
- **National data protection laws** vary significantly (China's PIPL, Brazil's LGPD, etc.)

Critical Infrastructure Protection:

- **NIS2 Directive** (EU) mandates cybersecurity for essential services — national implementation ongoing
- **CISA regulations** (U.S.) require critical infrastructure protection — updated in response to threats
- **National cybersecurity frameworks** vary by country and evolve continuously

⚠️ **These regulations are NOT static. Compliance requirements change, sometimes**

with minimal notice.

Get Berend Watchus's stories in your inbox

Join Medium for free to get updates from this writer.

Enter your email

Subscribe

This Article's Legal Status:

This analysis discusses:

- ✅ Publicly available academic research (arXiv papers, open conferences)
- ✅ Commercially available cryptographic standards (NIST PQC)
- ✅ Open-source software implementations (CrypTen, OpenSSL)
- ✅ Technical synthesis of unclassified information

This article does NOT:

- ❌ Provide export-controlled technical specifications
- ❌ Include ITAR-restricted weapon system designs
- ❌ Disclose classified information
- ❌ Offer guidance on evading export controls or regulatory requirements
- ❌ Constitute legal advice on compliance with any jurisdiction

⚠️ **LEGAL DISCLAIMER:**

This article does not constitute legal advice. Readers are **solely and entirely responsible** for:

1. **Determining current legal requirements** in their jurisdiction at the time of use
2. **Verifying export control status** of any technology before development, deployment, or transfer
3. **Ensuring compliance** with all applicable regulations, including but not limited to:

- Export controls (Wassenaar, EAR, ITAR, EU Dual-Use Regulation)
- Autonomous systems regulations (EU AI Act, national UAV laws)
- Privacy and data protection laws (GDPR, CCPA, national frameworks)
- Critical infrastructure protection requirements
- Sanctions and embargoes

1. **Consulting qualified legal counsel** before implementing any technologies discussed
2. **Monitoring regulatory changes** that may affect previously lawful activities

The legal status of dual-use technologies can change rapidly based on:

- Geopolitical developments
- National security assessments
- Technology evolution
- International incidents
- Legislative action
- Agency rulemaking

What is legal today may be restricted tomorrow. What is restricted in one jurisdiction may be unrestricted in another.

Ignorance of current regulations is not a defense. Due diligence is your responsibility.

This Work Focuses on Communication, Not Action:

This article analyzes:

- ✅ How secure LLM inference can be achieved (the communication mechanism)
- ✅ How hardware acceleration makes privacy-preserving coordination practical
- ✅ How existing research can be synthesized into deployable systems
- ✅ Strategic implications for IoT security and autonomous systems

This article does NOT provide:

- ❌ Operational procedures for specific missions
- ❌ Targeting methodologies or tactical doctrine
- ❌ Weapons integration or kinetic system designs
- ❌ Techniques to bypass security controls or harm individuals
- ❌ Methods to circumvent legal restrictions on encryption
- ❌ Procurement guidance for evading export controls

Ethical Considerations

CRITICAL: Ethical Frameworks are Context-Dependent and Evolving

⚠️ **Just as legal regulations change by jurisdiction and time, ethical frameworks for dual-use technology vary by cultural context, professional standards, and societal norms.**

What one society considers an acceptable use of surveillance technology, another

may view as a human rights violation. What one professional community considers responsible disclosure, another may view as reckless. What was ethically acceptable in 2020 may be ethically problematic in 2025 as societal understanding evolves.

⚠️ **READER RESPONSIBILITY: You are responsible for evaluating the ethical implications of any technology use within your specific cultural, professional, and societal context at the time of deployment.**

This article provides an ethical framework as of December 2025. Ethical standards evolve, and you must engage with current discourse in your community.

The Dual-Use Dilemma:

The alternative to developing privacy-preserving autonomous systems is NOT a world without autonomous systems — it is a world where only those willing to operate without privacy constraints deploy them.

Consider the banking analogy: When chip-and-PIN technology was introduced, banks didn't ask "Will criminals also benefit from secure transactions?" They asked "How do we protect legitimate customers from fraud?" The fact that criminals adapted to use the same secure infrastructure doesn't make the technology unethical — it makes it essential.

Consider the radio analogy: When digital encryption was added to P25 public safety radios, agencies didn't ask "Could organized crime groups purchase similar equipment?" They asked "How do we protect officer safety during operations?" The fact that the same radios are used today by both paramedics and organized crime doesn't make the technology dangerous — it makes coordinated emergency response possible.

Strong cryptography, secure inference, and quantum-safe hardware are defensive technologies that protect legitimate operations from:

- Surveillance and mass data collection
- Interception and eavesdropping

- Future quantum computing attacks (“Harvest Now, Decrypt Later”)
- Unauthorized access and manipulation
- Operational security compromises

Ethical Deployment Framework:

Building on my **Unified Model of Consciousness** research, autonomous systems with LLM-driven decision-making require built-in ethical constraints:

The Synthetic Insula Concept: An autonomous system must process both:

- External mission objectives (the “what”)
- Internal ethical boundaries (the “should”)

Just as the human insula integrates bodily states with cognitive processing, autonomous systems need architectural components that integrate mission parameters with ethical constraints. This isn’t about “making AI moral” — it’s about ensuring systems have feedback loops that can recognize constraint violations before action.

The Neglect Dilemma: (Explored in my UAV disaster response research)

Should an autonomous system prioritize:

- Adherence to programming (preventing misuse)?
- Or moral imperative to act (saving lives)?

This philosophical tension — action vs. inaction in critical scenarios — remains unresolved and requires human governance frameworks, not purely technical solutions.

Governance Requirements:

Technology neutrality does NOT mean governance neutrality. Deployment of secure

autonomous systems requires:

1. **Accountability Mechanisms**

- Audit trails for autonomous decisions
- Human oversight for high-consequence actions
- Clear chains of responsibility

1. **Transparency Standards**

- Open documentation of capabilities
- Public disclosure of deployment contexts
- Independent verification of safety claims

1. **Use Case Restrictions**

- Prohibited applications (e.g., autonomous lethal targeting without human control)
- Regulated applications (e.g., surveillance requiring judicial oversight)
- Permitted applications (e.g., disaster response, infrastructure monitoring)

1. **International Cooperation**

- Export control harmonization
- Shared safety standards
- Incident reporting frameworks

⚠️ **These governance requirements are aspirational and vary significantly by jurisdiction. Some nations have robust frameworks; others have none. Some professional communities have established ethical guidelines; others are still debating them.**

Your ethical responsibility extends beyond legal compliance to include:

- Engaging with professional ethical standards in your field
- Considering potential harms to affected communities
- Participating in governance discussions
- Refusing use cases you determine to be unethical, even if legal
- Advocating for better regulatory frameworks where needed

The Historical Lesson:

The inventors of foundational security technologies bear no responsibility for misuse:

- Whitfield Diffie and Martin Hellman (public-key cryptography) enabled secure internet commerce
- Ron Rivest, Adi Shamir, Leonard Adleman (RSA algorithm) made digital signatures possible
- Phil Zimmermann (PGP encryption) protected human rights activists and journalists
- The engineers at RSA Security (hardware tokens) secured banking for decades
- The developers of AES encryption (Rijndael) created the global standard for data protection
- Motorola, Hytera, and Kenwood engineers who built encrypted DMR radios enabled secure emergency communications

None of them are responsible for how encrypted channels are subsequently used. They built infrastructure that makes modern digital society, secure emergency response, and protected communications possible.

However, ethical responsibility for USE lies with deployers, operators, and governing authorities — not with researchers documenting capabilities.

The Uncomfortable Truth:

As you read this article, encrypted radios manufactured by legitimate companies are being used in:

- ✅ Local fire departments responding to structure fires
- ⚠️ Active armed conflicts in multiple regions globally
- ✅ Hospital emergency departments coordinating trauma care
- ⚠️ Organized crime operations in major cities worldwide
- ✅ Police departments protecting public safety
- ⚠️ Human trafficking networks coordinating illegal transportation

Should Motorola stop manufacturing encrypted radios because armed groups use them?

No. Because then local fire departments couldn't coordinate safely during hazmat incidents.

Should banks eliminate transaction encryption because money launderers use it?

No. Because then online banking would be vulnerable to interception.

Should we stop developing privacy-preserving AI systems because they have dual-use potential?

No. Because then only those willing to operate without privacy protections would deploy autonomous systems, and legitimate users (emergency responders, critical infrastructure operators, privacy-conscious civilians) would be left vulnerable.

⚠️ **But ethical responsibility requires asking:**

- Who will use this system?
- For what purpose?
- With what oversight?
- What harms might result?
- Are there less harmful alternatives?
- What governance mechanisms are in place?

These questions must be answered by deployers at the time of deployment, in their specific context, using current ethical frameworks in their community.

Responsible Disclosure:

All technical details in this analysis are derived from:

- Published academic papers (arXiv, peer-reviewed conferences)
- Open-source implementations (CrypTen, AirSim, OpenSSL)
- Publicly disclosed hardware specifications
- Established cryptographic standards (NIST PQC, AES-256, P25 specifications)
- Commercial product documentation
- FCC public databases and equipment authorizations

No classified information, proprietary details, or operational security vulnerabilities are disclosed or exploited in this work.

The Author's Position:

Technology advances regardless of individual preference. Banking apps will continue to use encryption whether or not we document how it works. Autonomous

systems will continue to develop whether or not we discuss their architecture. Digital radio will continue to offer encryption capabilities whether or not we acknowledge the dual-use nature.

Encrypted radios are being used right now by both paramedics and criminals. This doesn't make the technology unethical — it makes regulation, policy, and legal enforcement necessary while maintaining the infrastructure that protects legitimate users.

The question is not whether secure, LLM-driven autonomous swarms will exist, but whether their development includes:

- **Transparent research** enabling defensive countermeasures
- **Ethical frameworks** for deployment (as outlined in my Unified Model of Consciousness work)
- **Legal compliance** with export controls, privacy regulations, and autonomous systems governance
- **Accessible knowledge** that prevents monopolistic control
- **Public discourse** on governance and acceptable use

This article contributes to that discourse by making the technological convergence visible to researchers, policymakers, and security professionals who must navigate this landscape responsibly.

The Banking Parallel:

When EMV chip cards were introduced globally, banks didn't stop development because criminals would also benefit from more secure transactions. When banking apps replaced physical security tokens, financial institutions didn't halt progress because the same encryption could protect illicit transfers.

They developed the infrastructure, implemented it broadly, established regulations for acceptable use (KYC, AML compliance), and created law enforcement

mechanisms to address misuse — all while maintaining the fundamental security that protects billions of legitimate users.

The Radio Parallel:

When P25 Phase 2 encryption was standardized for public safety communications, agencies didn't refuse adoption because the same radios could be purchased commercially. When DMR became the global standard for digital land mobile radio, manufacturers didn't restrict sales because criminal organizations might acquire the equipment.

They knew — and accepted — that the same radios would be used simultaneously by emergency responders and armed groups in conflict zones. This didn't stop deployment. It informed policy, export controls, and regulatory frameworks.

They built the secure infrastructure, deployed it for legitimate emergency services, established frequency licensing and authorization procedures, and created regulatory frameworks to address misuse — all while maintaining the operational security that protects first responders and enables coordinated emergency response.

The Autonomous Systems Model:

This is the model for emerging autonomous and AI-driven systems:

1. **Develop the secure infrastructure** (PrivLLMSwarm's MPC framework + QSC-1 hardware acceleration)
2. **Establish governance frameworks** (international standards, use case restrictions, accountability mechanisms)
3. **Enable defensive capabilities** (allowing legitimate users to protect against surveillance and quantum threats)
4. **Address misuse through policy and legal enforcement** — not by leaving the infrastructure weak, secretive, or monopolistically controlled

Weak encryption doesn't stop criminals — it only harms legitimate users. The same principle applies to privacy-preserving autonomous systems.

Final Notes:

Legal Disclaimer: This article does not constitute legal advice. Readers are responsible for ensuring compliance with applicable export controls, autonomous systems regulations, privacy laws, and other legal requirements in their jurisdiction. Consult qualified legal counsel before implementing any technologies discussed.

Ethical Framework: The analysis of dual-use technologies in this article reflects the author's commitment to transparent research and responsible disclosure, consistent with academic freedom and the cybersecurity community's need for comprehensive threat modeling and defensive capability assessment.

Acknowledgment: This article builds upon the technical foundation established in "PrivLLMSwarm: Privacy-Preserving LLM-Driven UAV Swarms for Secure IoT Surveillance" (Ayana & Qiming, arXiv:2512.06747, 2025), extending that work with strategic synthesis, hardware integration proposals, legal analysis, and ethical considerations not present in the original publication.

Author: Berend Watchus
Date: December 10, 2025

— — — — — — — — — — — — — — — — —

also read:

HAWK TUAH: Why This Drone Mothership (U-Hawk) Matters

Beyond the Battlefield

HAWK TUAH: Why This Drone Mothership (U-Hawk) Matters Beyond the Battlefield Author: Berend Watchus October 22, 2025...

systemweakness.com

🔐 NATO Counter-Drone Operations in Belgium: Deep Technical Analysis for Cyber Security...

Author: Berend Watchus

systemweakness.com

P vs NP in the Real World: How Drones Navigate the Traveling Salesman Problem Through Universal...

P vs NP in the Real World: How Drones Navigate the Traveling Salesman Problem Through Universal Compression Principles...

systemweakness.com

FBI Seeks (face recogn) AI-Powered Drone Surveillance: Executive Brief & Comprehensive Analysis

FBI Seeks (face recogn) AI-Powered Drone Surveillance: Executive Brief & Comprehensive Analysis Author: Berend Watchus...

systemweakness.com

https://www.scribd.com/document/954333936<<

Drones | Drone Swarm | Infosec | AI | Cyber Security Awareness

Follow

Written by Berend Watchus

26 followers · 45 following

Independent Researcher (Non-Profit) exploring AI & interdisciplinary solutions. Former futurist event organizer & care innovation lab volunteer

No responses yet

Write a response

What are your thoughts?

More from Berend Watchus

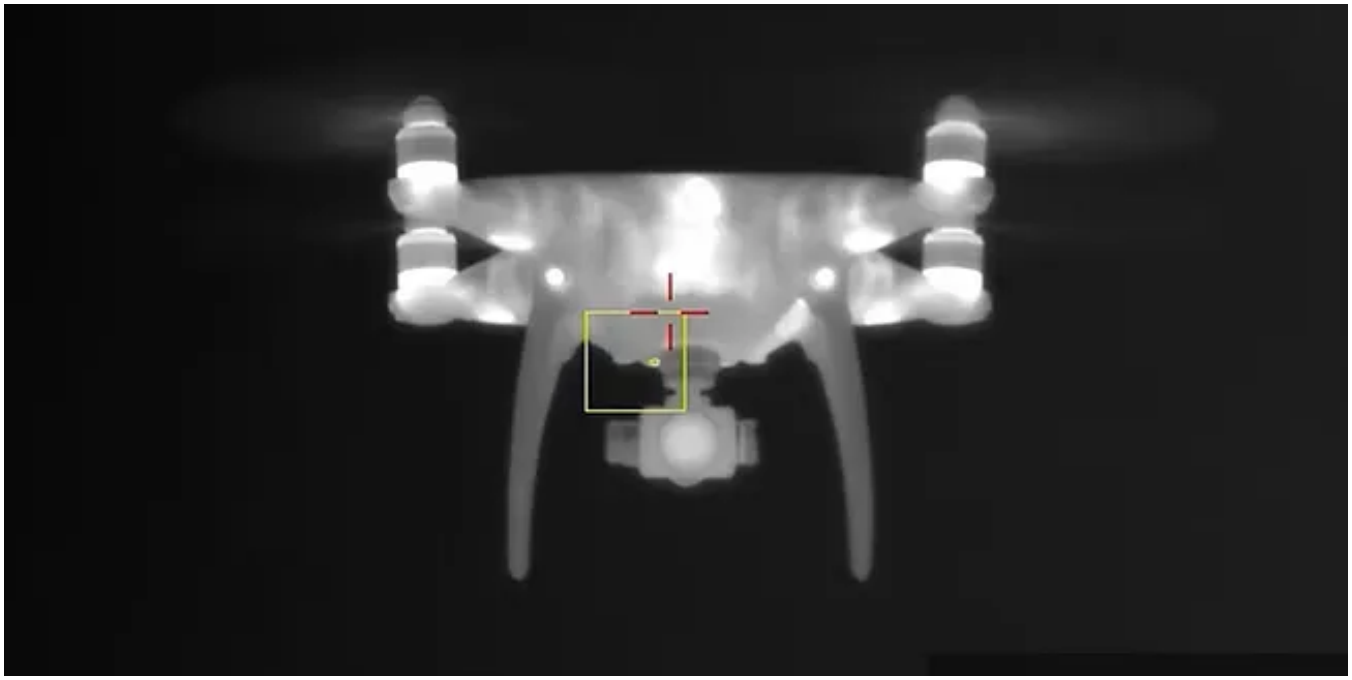

In System Weakness by Berend Watchus

🔐 NATO Counter-Drone Operations in Belgium: Deep Technical Analysis for Cyber Security...

Author: Berend Watchus

Nov 12

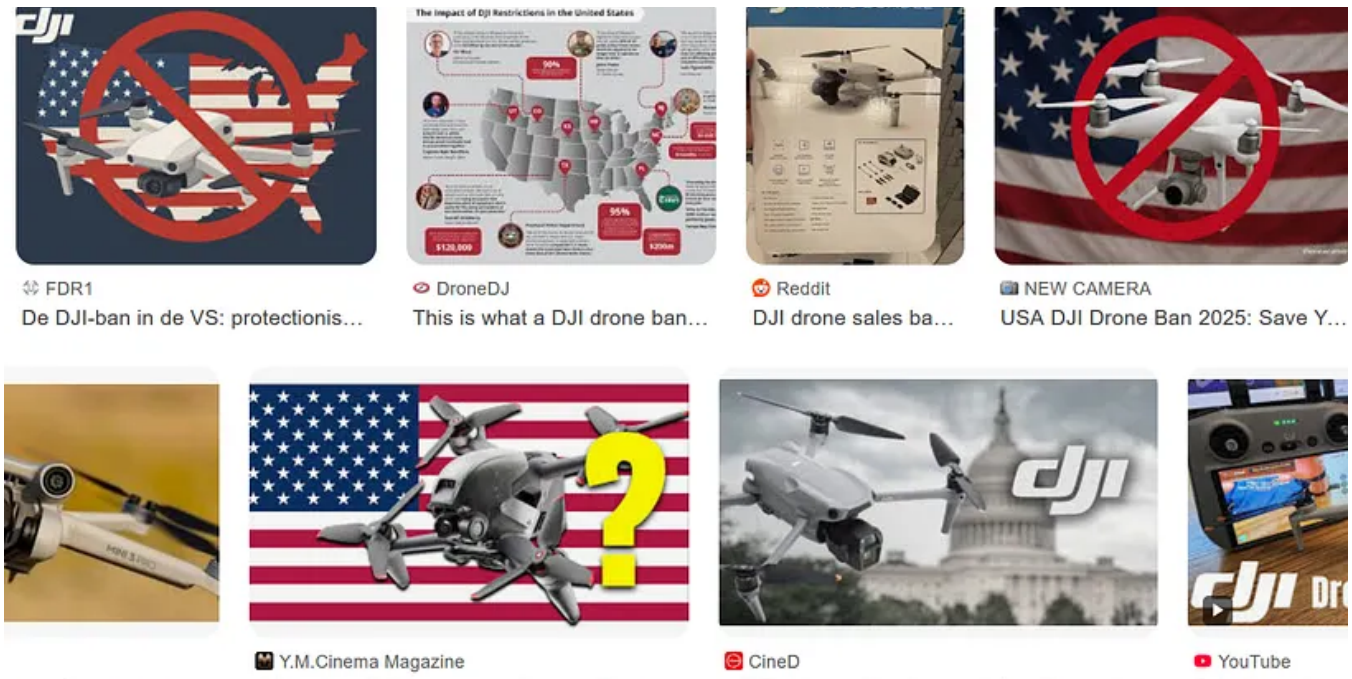

In System Weakness by Berend Watchus

The DJI Drone Ban: Comprehensive 2025 Update for Cybersecurity Professionals

Author: Berend Watchus includes free download in pdf

Nov 24

In System Weakness by Berend Watchus

Linguistic Trojan Horse: Why LLM Steganography Just Broke AI Safety

Author: Berend Watchus October 24, 2025 Blog article

Oct 24 1

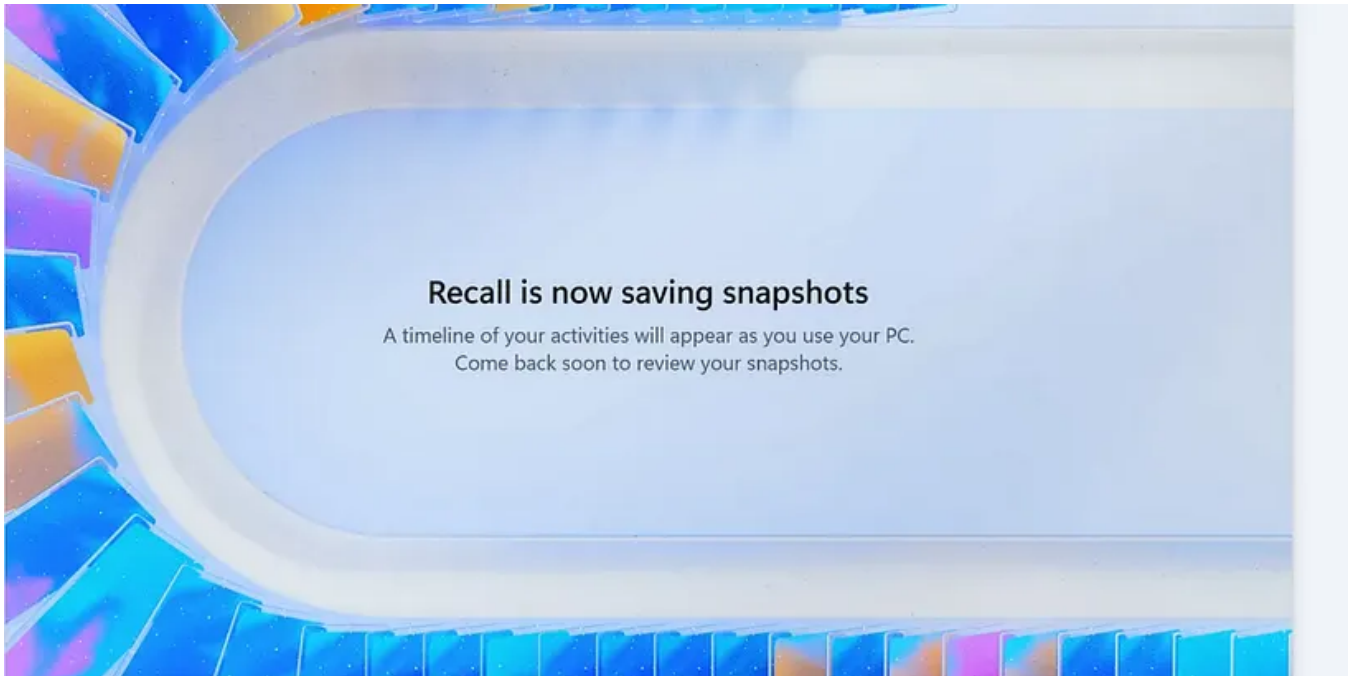
Recall is now saving snapshots
A timeline of your activities will appear as you use your PC.
Come back soon to review your snapshots.

In System Weakness by Berend Watchus

Windows Recall: The VBS Enclave Paradox and the Evolution of Endpoint Espionage

Author: Berend Watchus November 11, 2025 Blog article

Nov 11 1

See all from Berend Watchus

Recommended from Medium

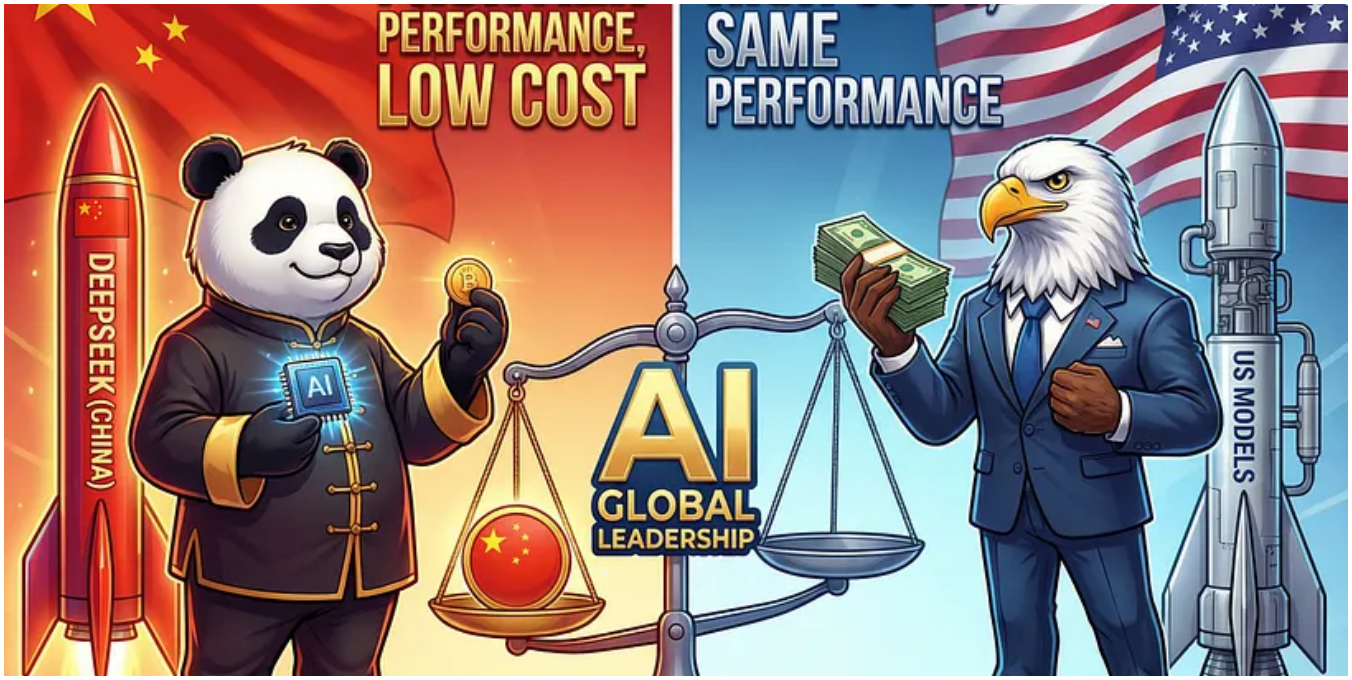

Ignacio de Gregorio

The New Chinese Humiliation

When cheaper is an understatement

4d ago 1K 43

Lucy Foulkes

Welcome to my body

After twenty years of pain and repeated medical dead ends, a stranger sent me a message on Instagram. It led to a diagnosis all the doctors...

In Artificial Corner by The PyCoach

The Best AI Tools for 2026

If you're going to learn a new AI tool, make sure it's one of these

Dec 1 1.3K 38

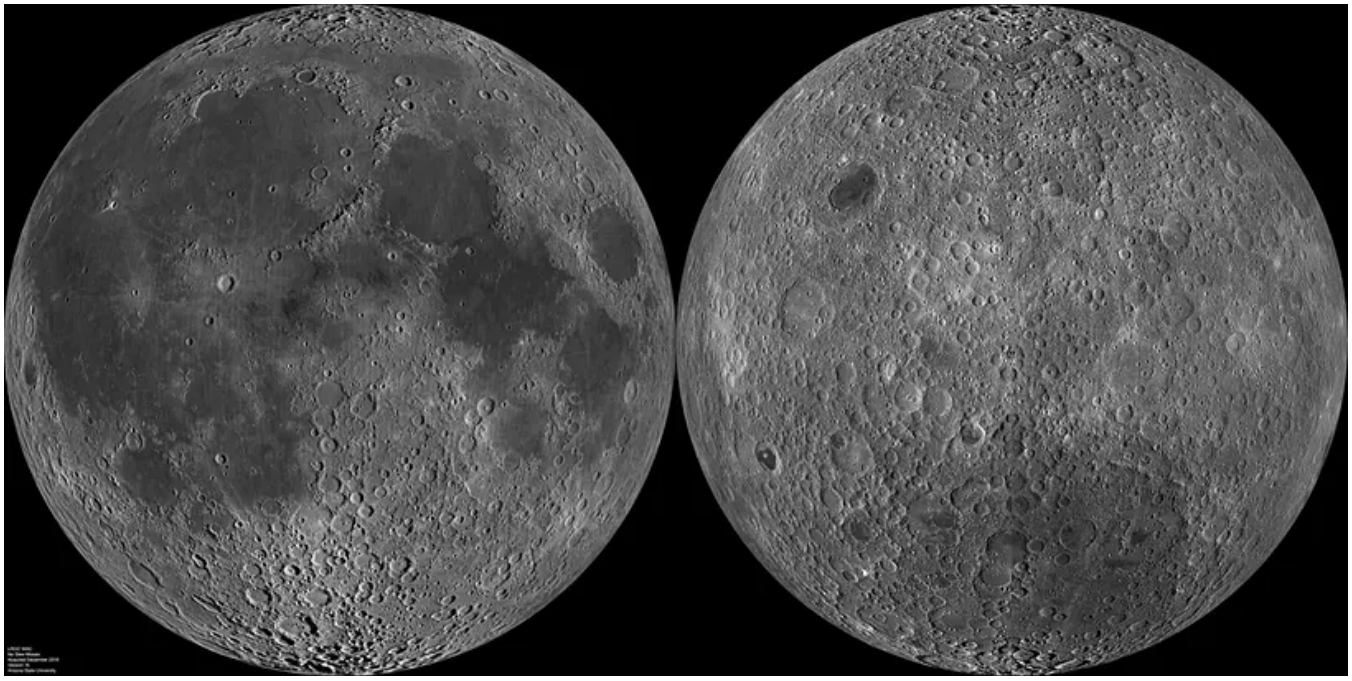

In Starts With A Bang! by Ethan Siegel

The Moon’s two faces don’t match, and we think we know why

The far side of the Moon is incredibly different from the Earth-facing side. 66 years later, we know why the Moon’s faces are not alike.

Dec 3 1.1K 14

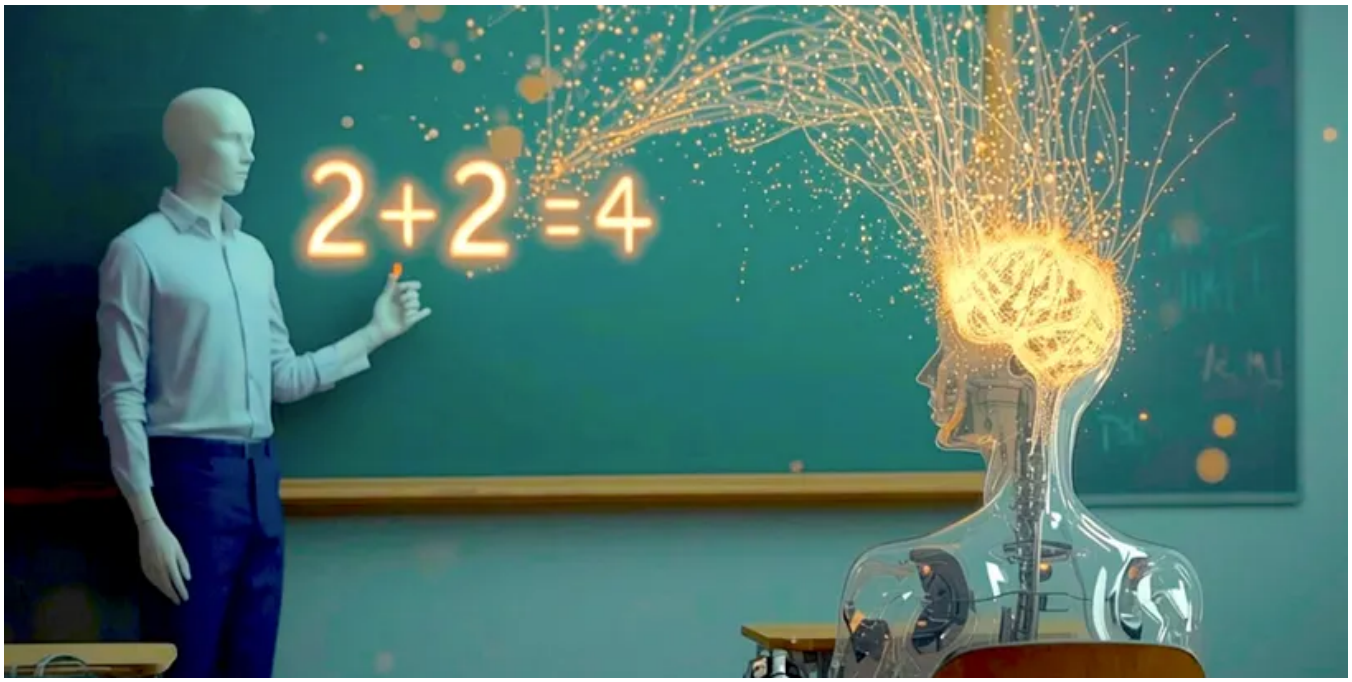

In AI Advances by Nikhil Anand

I wasted years confused about how LLMs learn before learning this.

I wish someone had told me this when I started.

Nov 30 620 13

In Predict by Rafe Brena, Ph.D.

Apple Smart Glasses Will Defeat Meta's ones at their Own Game

What is Ambient AI, anyway?

5d ago 404 13

See more recommendations